

Global Simulation of Plasma Microturbulence at the Petascale and Beyond

William M. Tang

**Princeton University, Plasma Physics Laboratory
Princeton, NJ USA**

Early Science Program Principal Investigators Meeting
“A Presentation of Mira’s First Science”

**Argonne Leadership Computing Facility
Argonne National Laboratory
Argonne, IL**

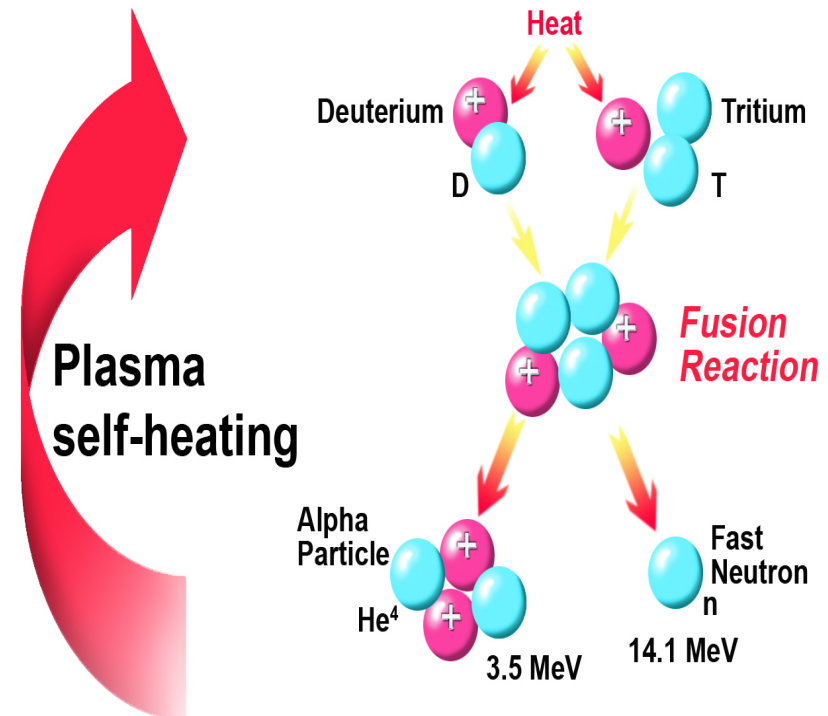
May 16, 2013

Fusion: an Attractive Energy Source

- **Abundant fuel, available to all nations**
 - Deuterium and lithium easily available for millions of years
- **Environmental advantages**
 - No carbon emissions, short-lived radioactivity
- **Cannot “blow up or melt down,” resistant to terrorist attack**
 - Less than minute’s worth of fuel in chamber
- **Low risk of nuclear materials proliferation**
 - No fissile materials required
- **Compact relative to solar, wind and biomass**
 - Modest land usage
- **Not subject to daily, seasonal or regional weather variation & no requirement for local CO₂ sequestration**
 - Not limited by need for large-scale energy storage nor for long-distance energy transmission
- ***Fusion is complementary to other attractive energy sources***

Fusion Energy: *Burning plasmas are self-heated and self-organized systems*

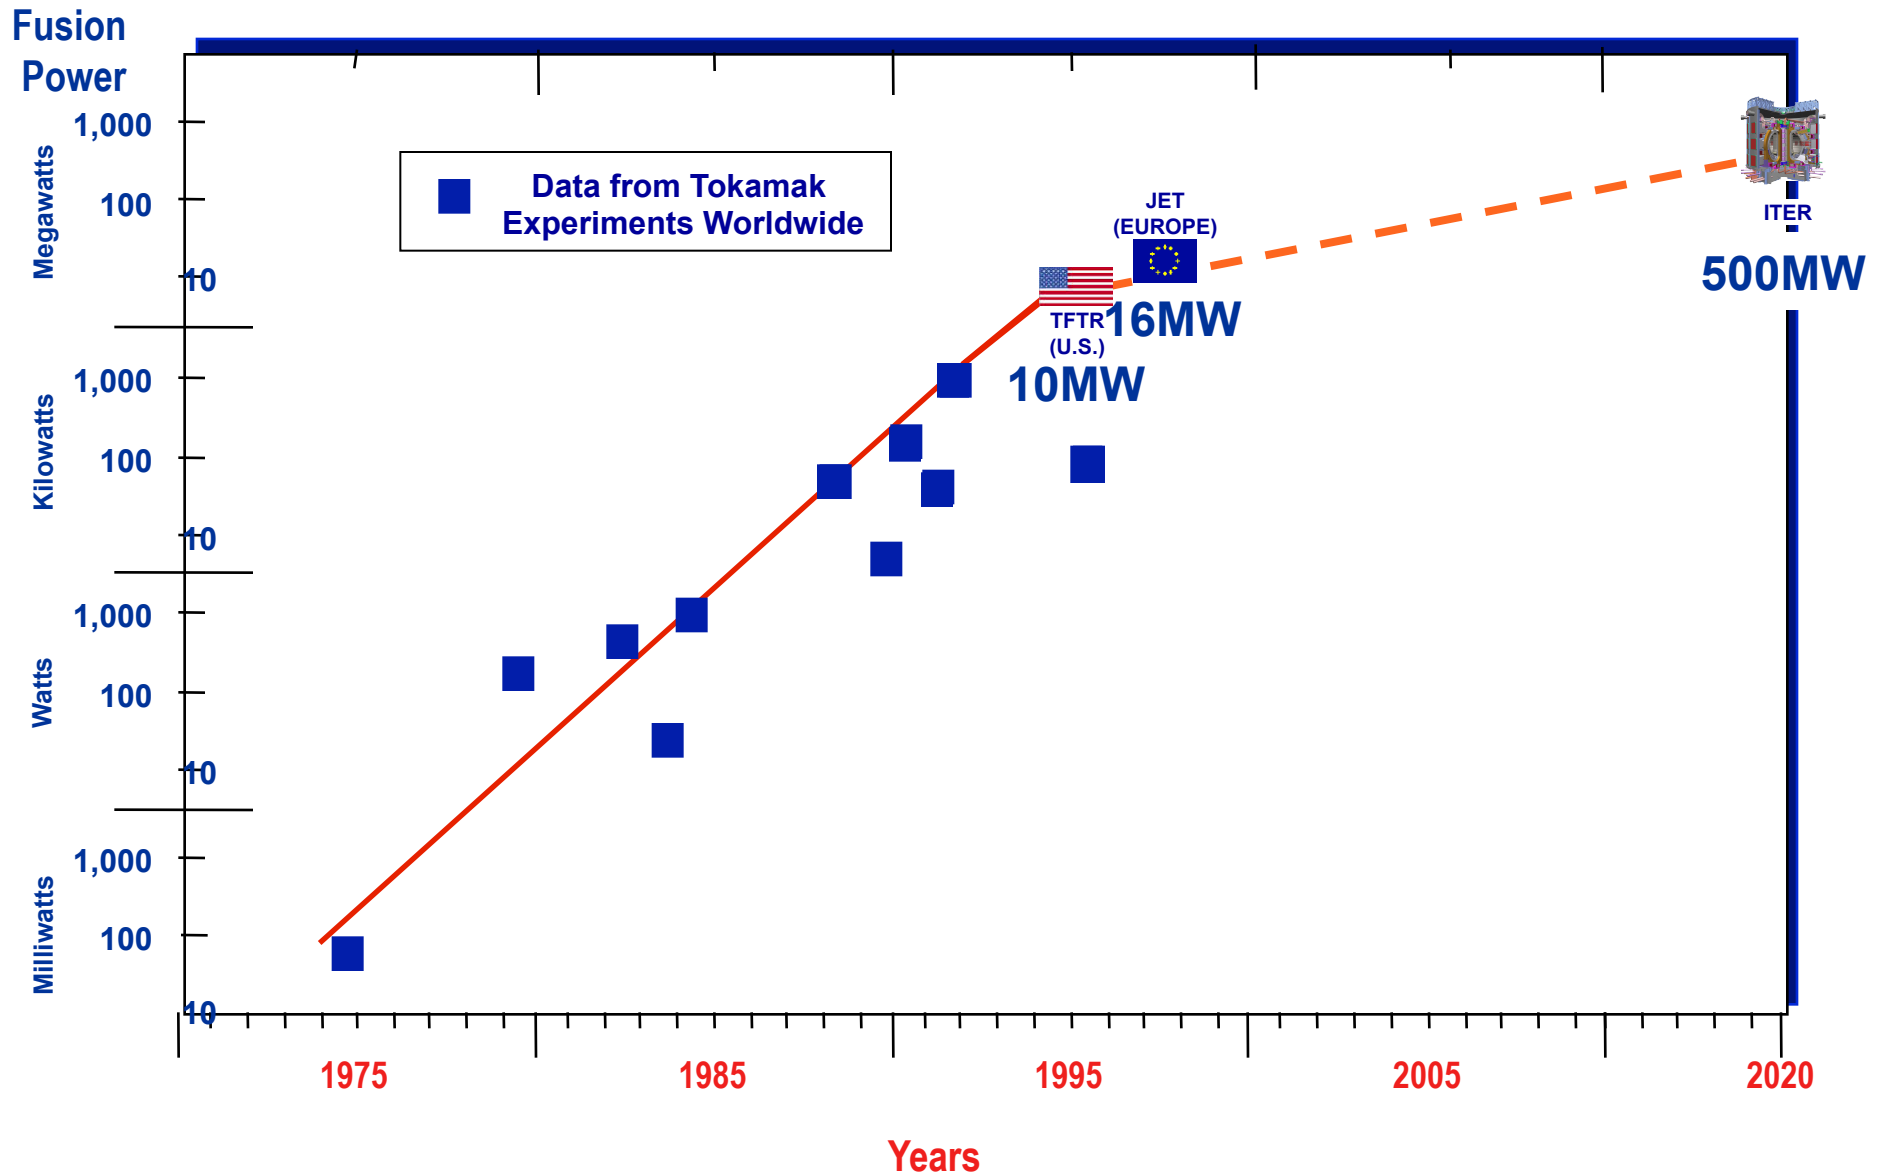
Deuterium-Tritium Fusion Reaction



***Energy Multiplication
About 450:1***

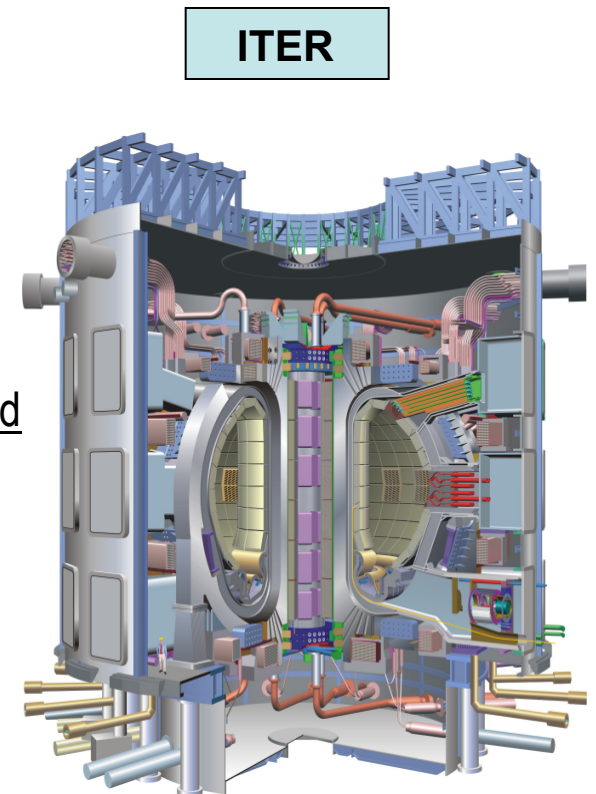


Progress in Magnetic Fusion Energy (MFE) Research

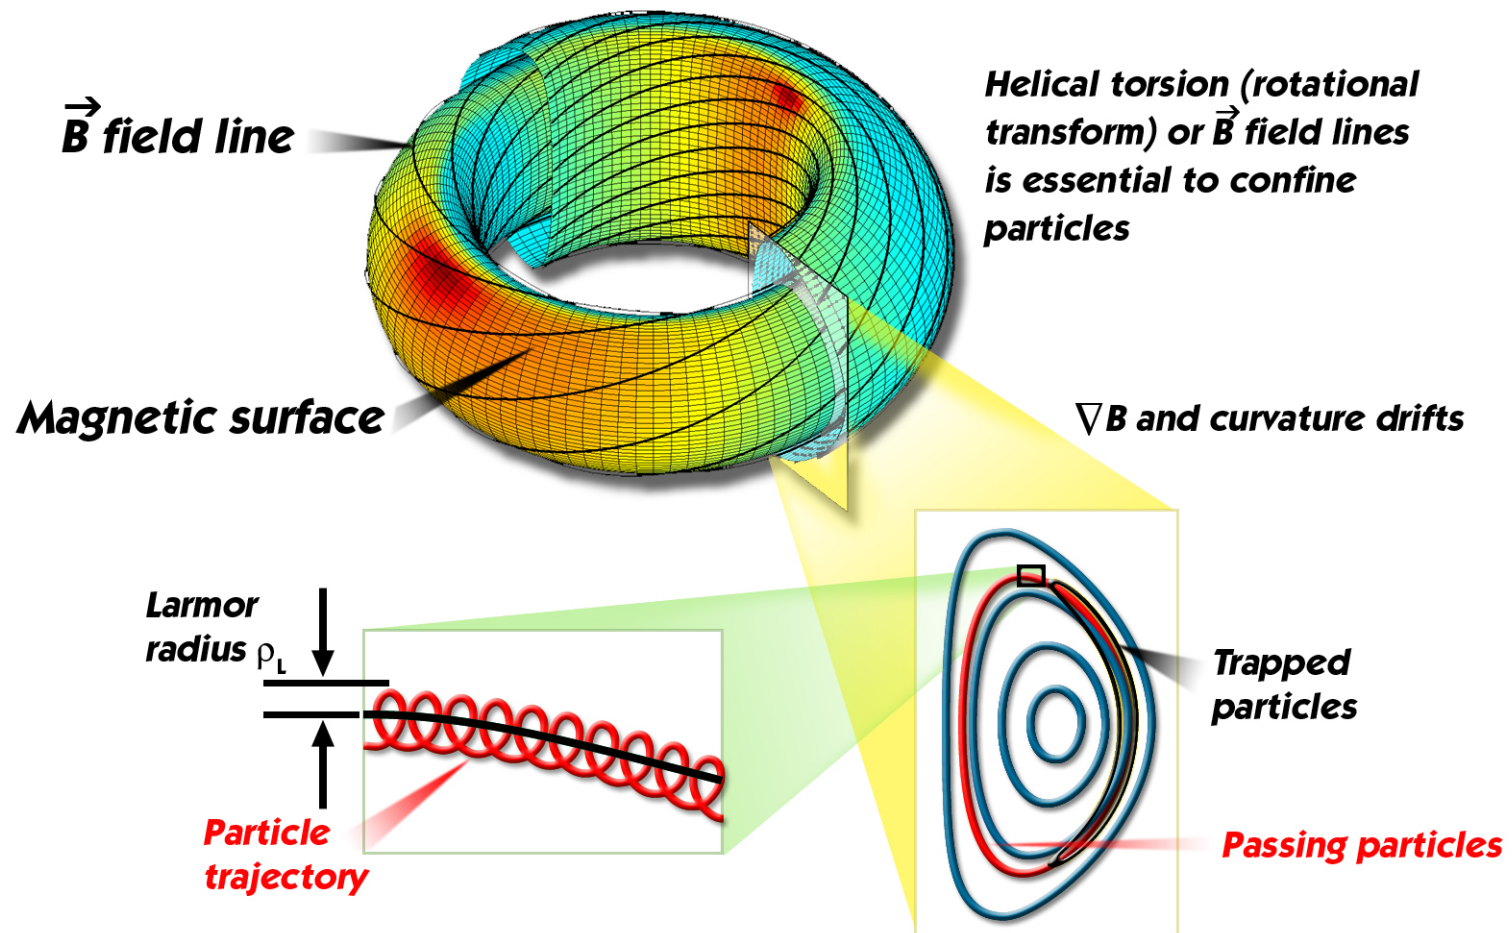


ITER Goal: *Demonstration of the Scientific and Technological Feasibility of Fusion Power*

- **ITER** is an ~\$20B facility located in France & involving 7 governments representing over half of world's population
→ **dramatic next-step for Magnetic Fusion Energy (MFE) producing a sustained burning plasma**
 - Today: 10 MW(th) for 1 second with gain ~ 1
 - ITER: 500 MW(th) for >400 seconds with gain >10
- **“DEMO”** will be demonstration fusion reactor after ITER
 - 2500 MW(th) continuous with gain >25 , in a device of similar size and field as ITER
- Ongoing R&D programs worldwide [experiments, theory, **computation**, and technology] essential to provide growing knowledge base for ITER operation targeted for ~ 2020
- **Realistic HPC-enabled simulations required to cost-effectively plan, “steer,” & harvest key information from expensive (~\$1M/long-pulse) ITER shots**



Magnetically confined plasmas in a tokamak are complex and require HPC analysis



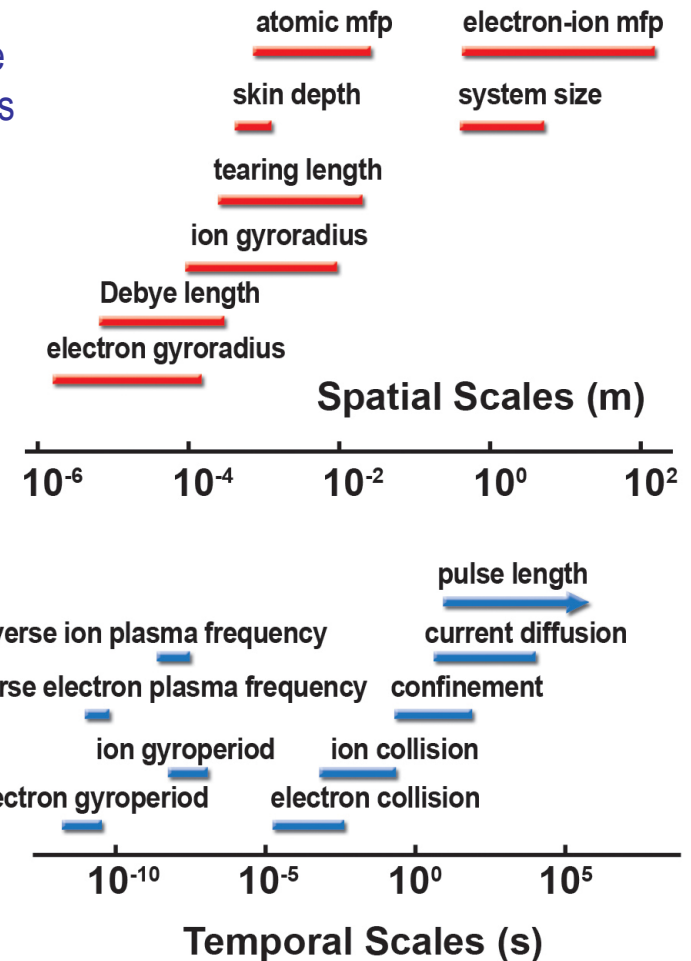
Modern HPC-enabled simulations open opportunities for “transformational” science to accelerate understanding in fusion energy research

Though equations are well-known (Boltzmann-Maxwell), the problem is a physics grand challenge

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f + \frac{q}{m} [\mathbf{E} + \mathbf{v} \times \mathbf{B}] \cdot \nabla_{\mathbf{v}} f = C(f) + S(f)$$

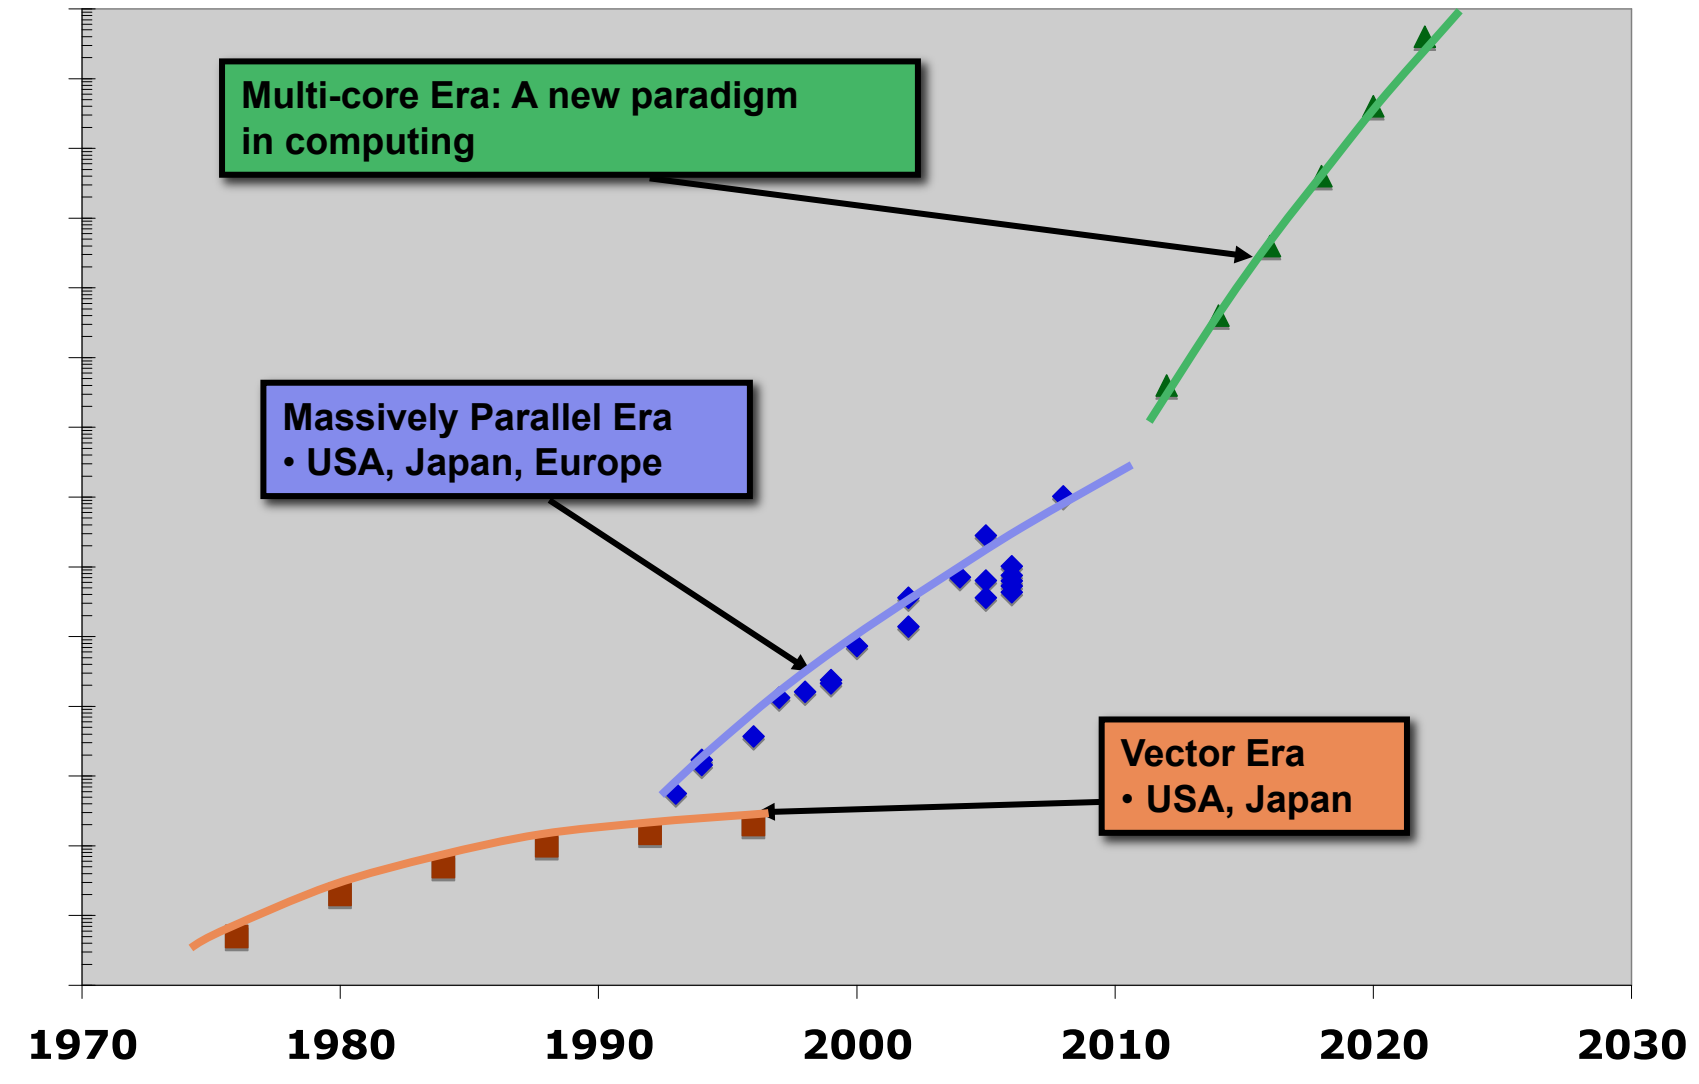
convection in space convection in velocity space collisional relaxation particle sources

- Seven dimensional equation of motion in phase space, $f(\mathbf{x}, \mathbf{v}, t)$ for each species and 2 coupled vector fields
- Extreme range of time scales – wall equilibration/electron cyclotron $O(10^{14})$
- Wide range of spatial scales – machine radius/electron gyroradius $O(10^4)$
- Extreme anisotropy – mean free path in magnetic field parallel/perpendicular $O(10^8)$
- Intrinsic nonlinearity (e.g. plasma distributions generate significant E and B fields through Maxwell’s equations)
- Sensitivity to geometric details
- *Advanced simulations required to address grand challenges in plasma science*



FES Needs to be Prepared to Exploit *Local Concurrency* to Take Advantage of Most Powerful Supercomputing Systems in 21st Century

(e.g., U.S.'s Blue-Gene-Q & Titan, Japan's Fujitsu-K, China's Tianhe-1A,)



Particle Simulation of the Boltzmann-Maxwell System

- The Boltzmann equation (Nonlinear PDE in Lagrangian coordinates):

$$\frac{dF}{dt} = \frac{\partial F}{\partial t} + \mathbf{v} \cdot \frac{\partial F}{\partial \mathbf{x}} + \left(\mathbf{E} + \frac{1}{c} \mathbf{v} \times \mathbf{B} \right) \cdot \frac{\partial F}{\partial \mathbf{v}} = C(F).$$

- “Particle Pushing” (Linear ODE’s)

$$\frac{d\mathbf{x}_j}{dt} = \mathbf{v}_j, \quad \frac{d\mathbf{v}_j}{dt} = \frac{q}{m} \left(\mathbf{E} + \frac{1}{c} \mathbf{v}_j \times \mathbf{B} \right)_{\mathbf{x}_j}.$$

- Klimontovich-Dupree representation,

$$F = \sum_{j=1}^N \delta(\mathbf{x} - \mathbf{x}_j) \delta(\mathbf{v} - \mathbf{v}_j),$$

- Poisson’s Equation: (Linear PDE in Eulerian coordinates (lab frame))

$$\nabla^2 \phi = -4\pi \sum_{\alpha} q_{\alpha} \sum_{j=1}^N \delta(\mathbf{x} - \mathbf{x}_{\alpha j})$$

- Ampere’s Law and Faraday’s Law [Linear PDE’s in Eulerian coordinates (lab frame)]

Particle-in-Cell Simulations

- Early attempts [*Buneman (1959)*; *Dawson (1962)*]
- Finite-Size Particles and Particle-in-Cell Simulation [*Dawson et al. (1968)* and *Birdsall et al. (1968)*]

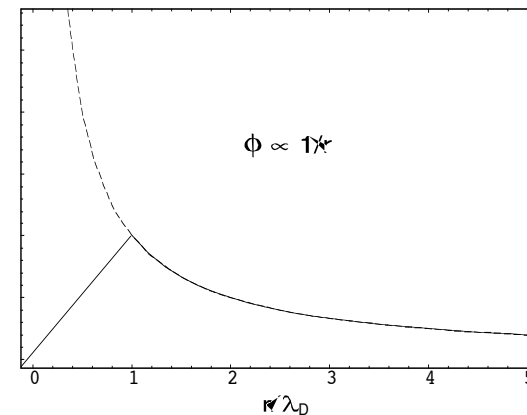
- Coulomb potential is modified for a finite size particle due to Debye shielding

- no need to satisfy $1/(n \lambda_D^3) \ll 1$

- Number of calculations for N particles

- N^2 for direct interactions and N for PIC

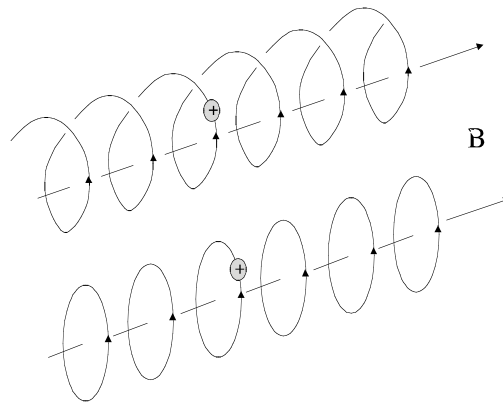
- Collisions are treated as sub-grid phenomena via Monte-Carlo methods [*Shanny, Dawson & Greene (1976)*]



Gyrokinetic Particle Simulation

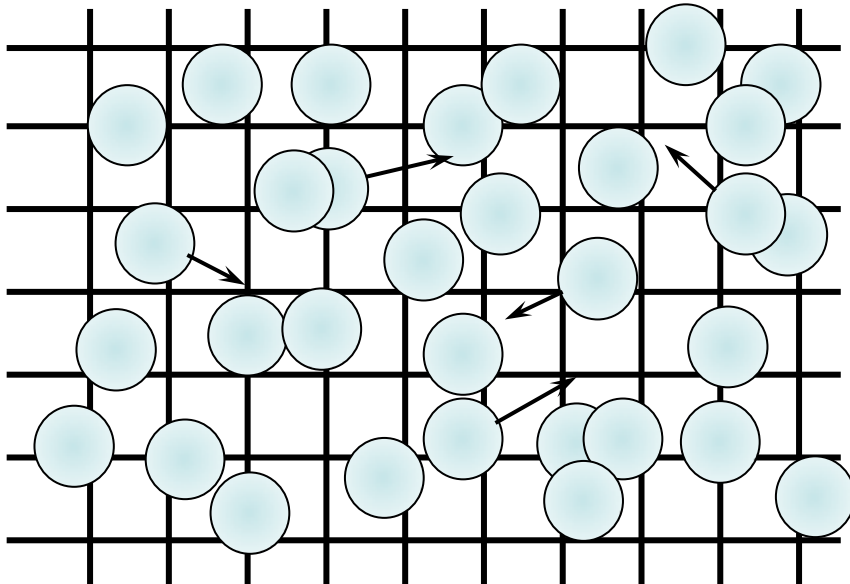
Ref. [W. W. Lee, *PF* ('83); *JCP* ('87)]

- Gyrophase-averaged Vlasov-Maxwell equations for low frequency microinstabilities.
- Spiral motion of a charged particle is modified as a rotating charged ring subject to guiding center electric and magnetic drift motion as well as parallel acceleration -- *speeds up computations by 3 to 6 orders of magnitude in time steps and 2 to 3 orders in spatial resolution*



Basic Particle-in-Cell Method

- Charged particles sample distribution function
- Interactions occur on a grid with the forces determined by gradient of electrostatic potential (calculated from deposited charges)
- Grid resolution dictated by Debye length (“finite-sized” particles) up to gyro-radius scale



Specific PIC Operations:

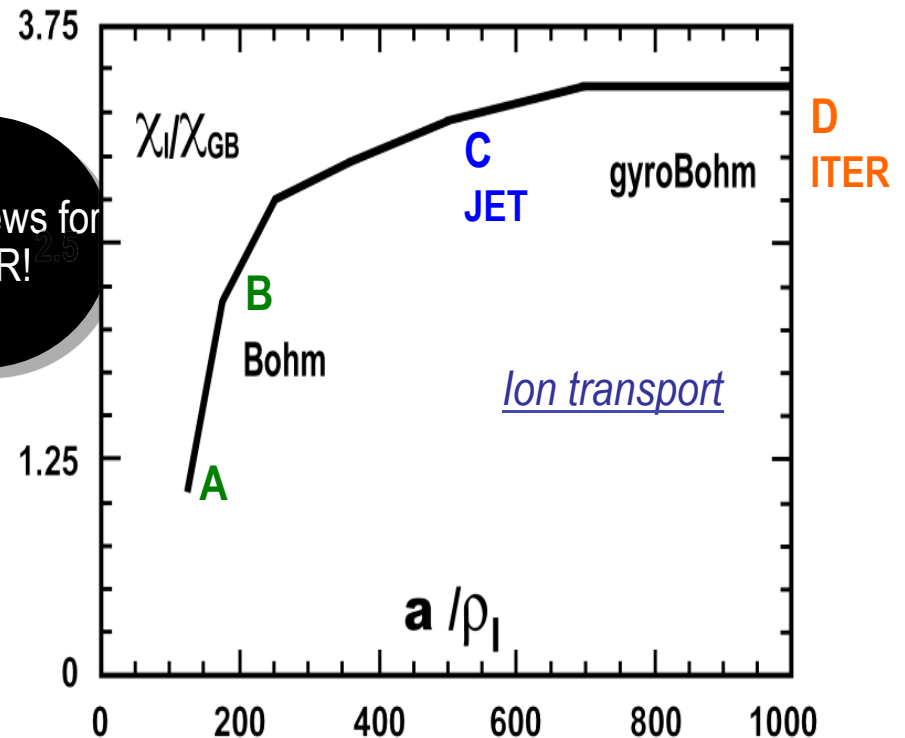
- “**SCATTER**”, or deposit, charges as “nearest neighbors” on the grid
- Solve Poisson Equation for potential
- “**GATHER**” forces (gradient of potential) on each particle
- Move particles (**PUSH**)
- Repeat...

Microturbulence in Fusion Plasmas – Mission Importance: Fusion reactor size & cost determined by balance between loss processes & self-heating rates

- “**Scientific Discovery**” - *Transition to favorable scaling of confinement produced in simulations for ITER-size plasmas*
 - $a/\rho_i = 400$ (JET, largest present lab experiment) through
 - $a/\rho_i = 1000$ (ITER, ignition experiment)
- *Multi-TF simulations* using GTC global PIC code [Z. Lin, et al, 2002) deployed a billion particles, 125M spatial grid points; 7000 time steps @ NERSC → *1st ITER-scale simulation with ion gyroradius resolution*
- **Understanding** physics of plasma size scaling demands *much greater computational resources* + improved algorithms [radial domain decomposition, hybrid (MPI+Open MP) language, ..] & modern diagnostics

→ ESP & INCITE GTC-P Projects on BG-Q @ ALCF

Good news for ITER!



→ **Excellent Scalability of Global PIC Codes on modern HPC platforms enables much greater resolution/physics fidelity to improve understanding**

→ **BUT - further improvements for efficient usage of current LCF's demands code re-write featuring modern CS/AM methods addressing locality, low-memory-per-core,**

Performance Speed Up Comparison Results (IBM BG-Q vs. BG-P)

M0180 ppc =100	Our test	ANL	IBM
Speed up per node (Q/P ratio)	10.7	10.7	11.2

Speed up per node ***comparison with ALCF and IBM results*** for “M0180” problem size (i.e., 180 grid-points in radial direction) using GTC-P code

→ Test Case for phase-space resolution with particles/cell (ppc) =100 for 100 time-steps

→ “Time to Solution” improvement from **BG-Q hardware**
• 4X (core) 2X (frequency) 2X (SIMD) = 16 (“theoretical”)

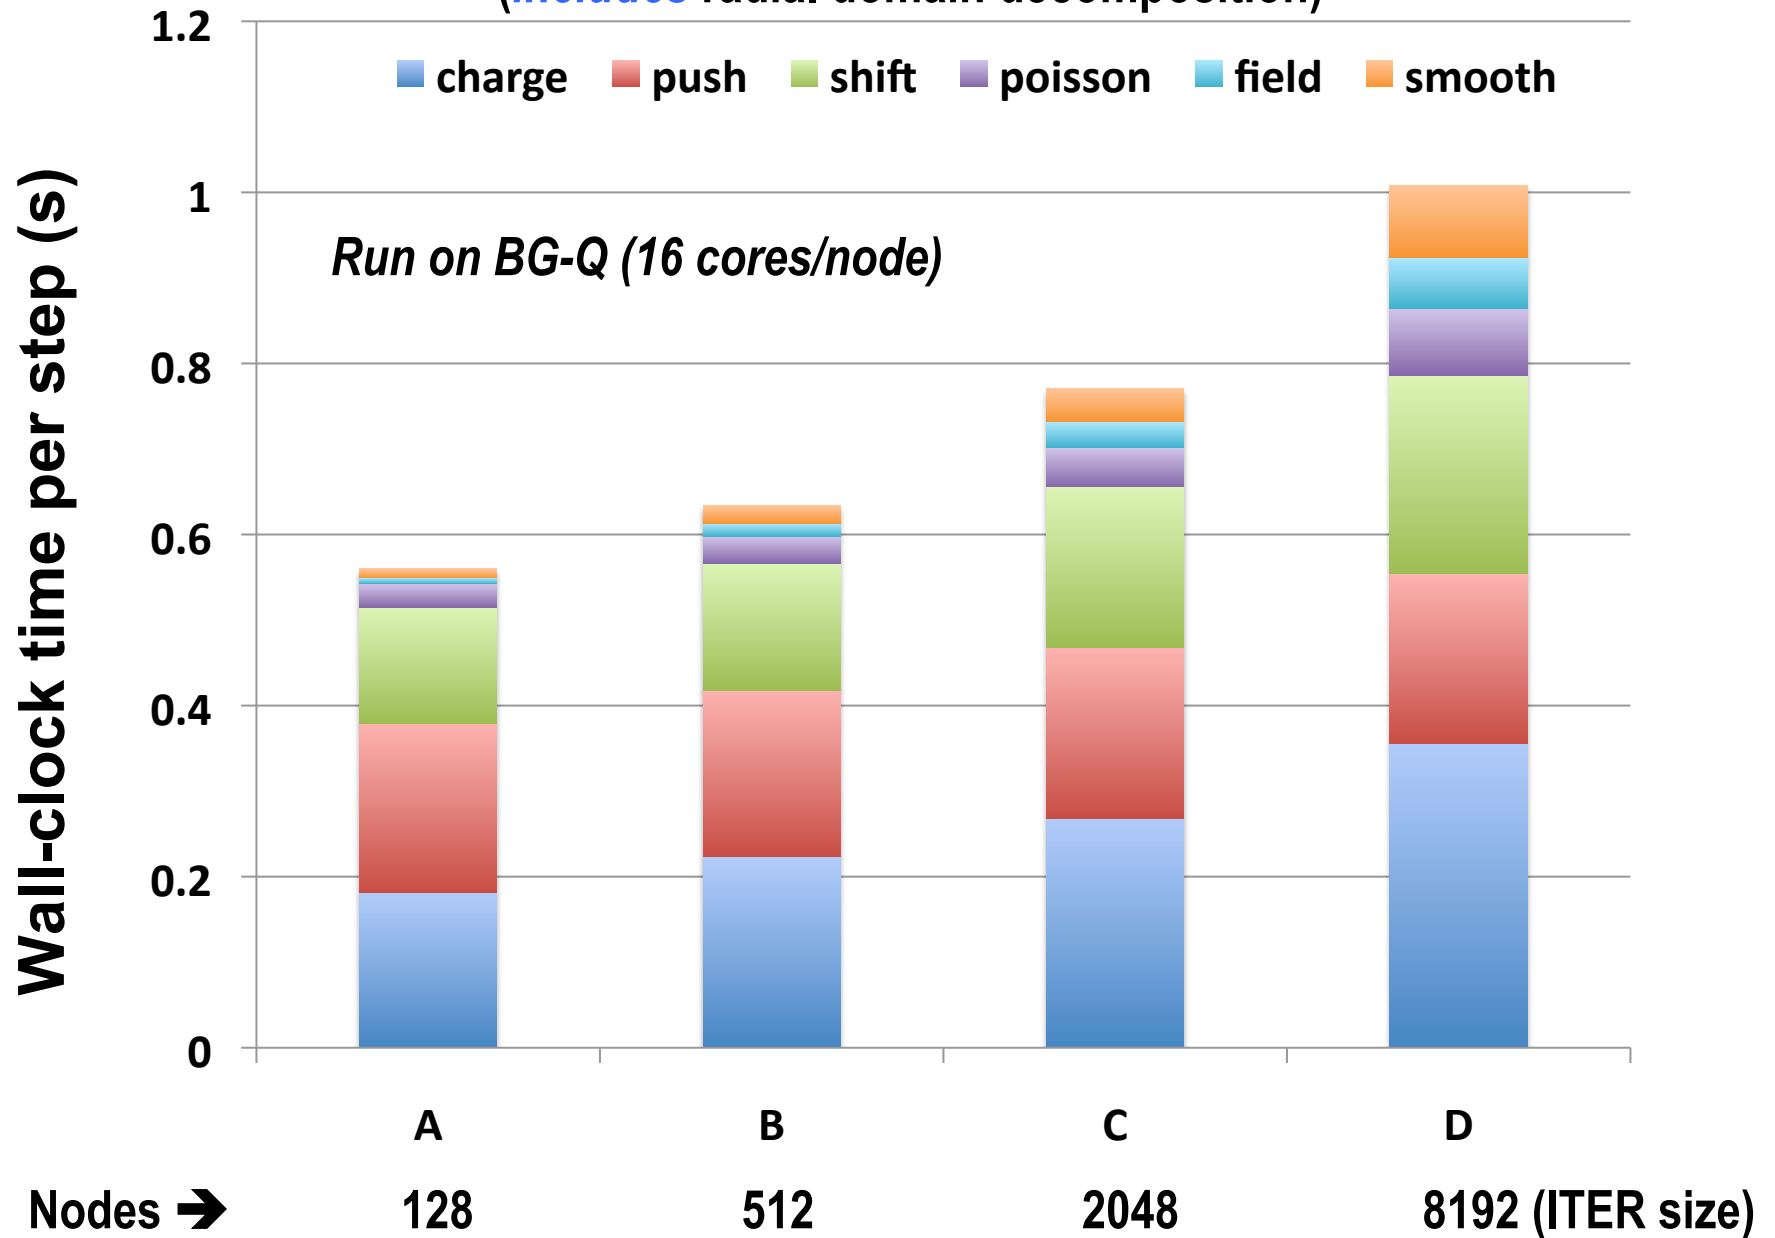
Features of new “GTCP-C” Code

- “GTCP-C” code based on a greatly optimized version of the C-version of the original GTC code (introduced at SC 2011)
 - “C” (instead of usual Fortran) to best incorporate CS community advances in multi-threading for low-memory-per-core systems
- Key additional level of **domain decomposition introduced into radial dimension**
→ essential for efficiently carrying out simulations on large-scale plasmas such as ITER
 - Alleviates grid memory requirement issue for large size plasma simulation
 - Improves locality
- Multiple levels of parallelism
 - 2D domain decomposition (toroidal and radial dimensions)
 - Particle decomposition in each domain
 - Multi-threaded, shared memory parallelism implemented with loop-level OpenMP directives
- Improvements over GTCP-FORTRAN code
 - Remove PETSc library for carrying out Poisson field solve
 - Significantly improves code portability to various LCF’s
 - Introduces loop level parallelism in the Poisson solve

→ Overall Software Improvement gives another 50% gain in “Time to Solution”

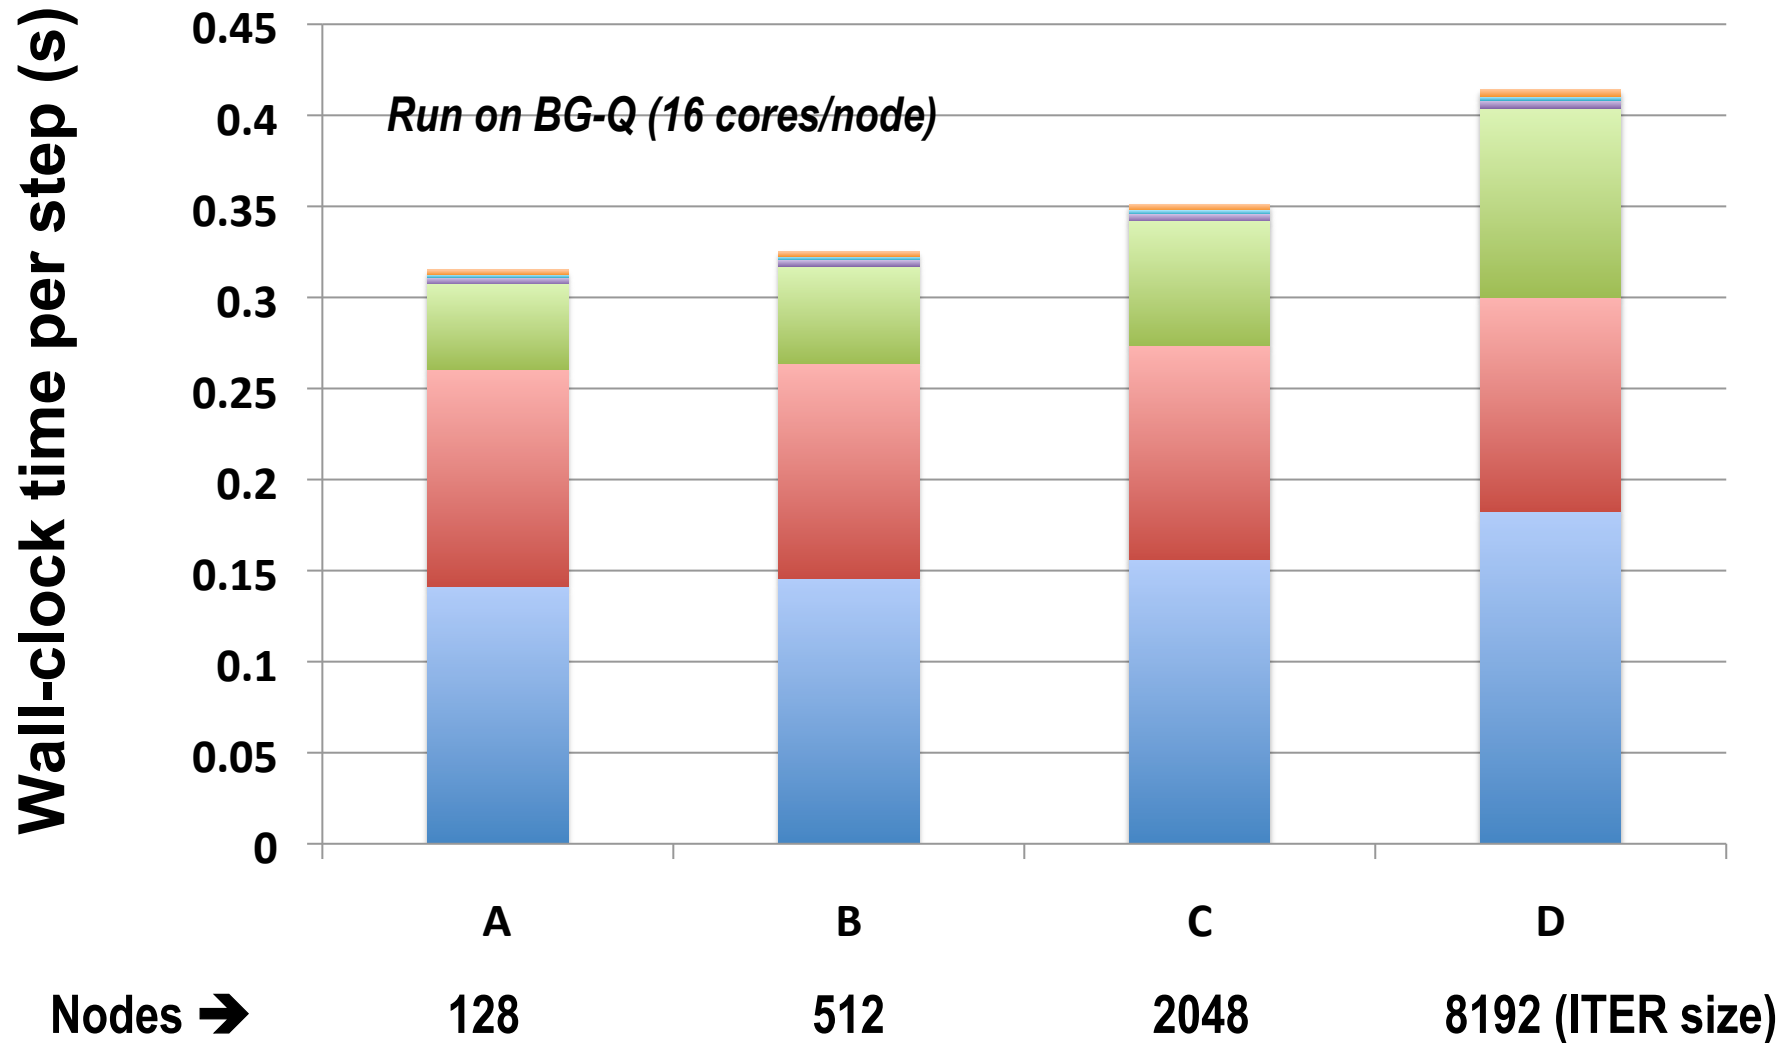
Reference Fortran Version of GTC-P

(*includes* radial domain decomposition)



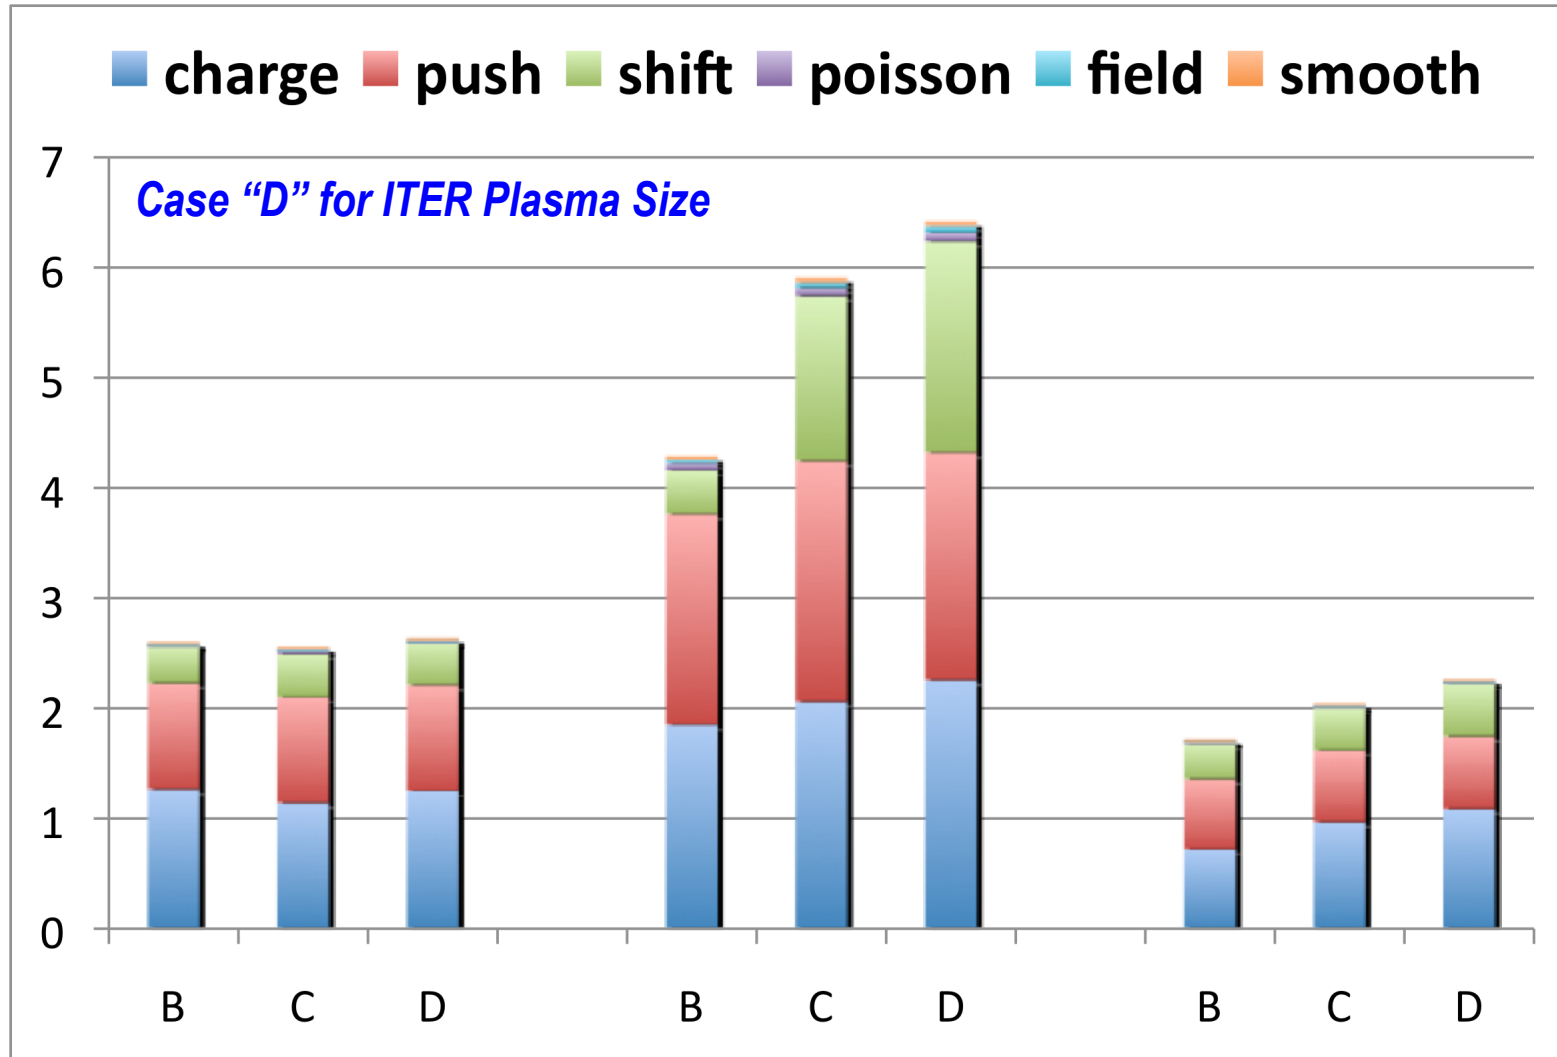
Optimized C Version of GTC-P [“GTC-P C”]

charge push shift poisson field smooth



Comparative Performance for PIC Operations

Wall-clock time per step (s)



Mira 1MPI/

64 OpenMP

[16 cores w/each launching 4 threads]

Hopper 4MPI/

6 OpenMP

[Intel design -- avoid NUMA effects]

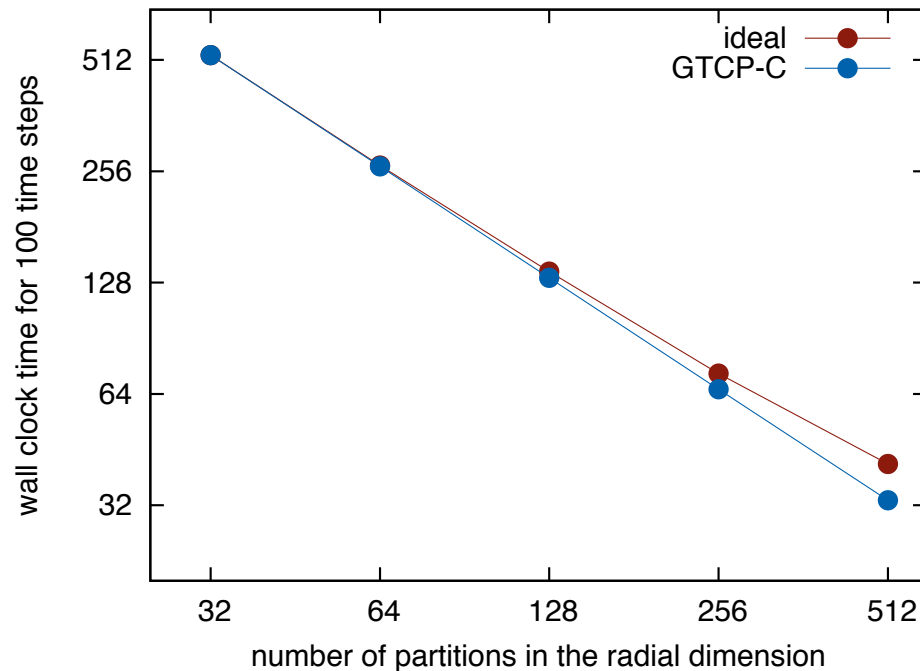
Edison 2MPI/

8 OpenMP

[avoid NUMA effects]

Strong Scaling Study of GTCP-C on Mira

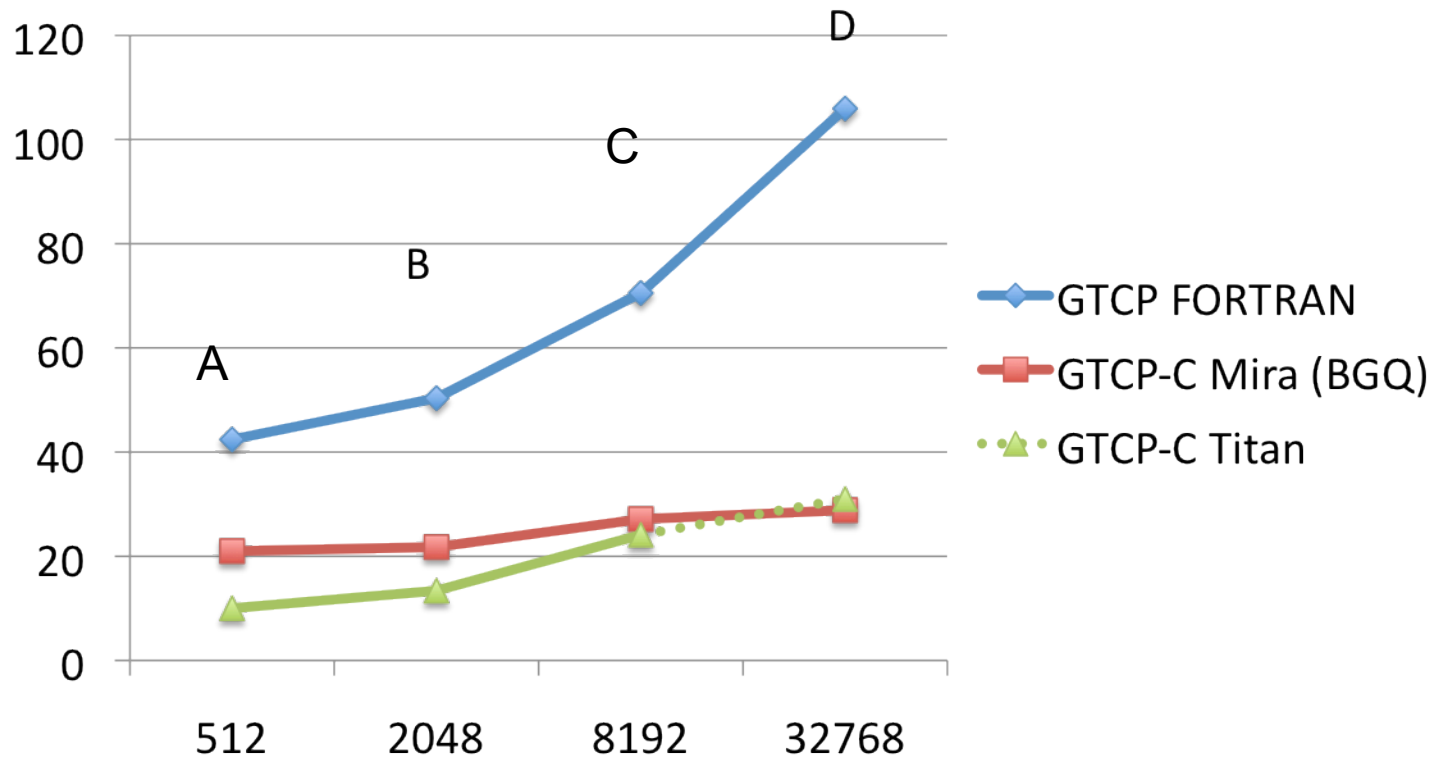
D (ITER-scale) Problem for 100 Time Steps



Radial partitions	Time on Mira	Ideal	"Eff" efficiency
32	527.5	527.5	100%
64	265.1	263.8	99%
128	137.1	131.9	96%
256	72.6	65.9	91%
512	41.4	33	80%

- 64-way toroidal partitioning on all numerical experiments
- BG/Q (Mira) system --≥ **use 4 processes/node, 16 threads/process**

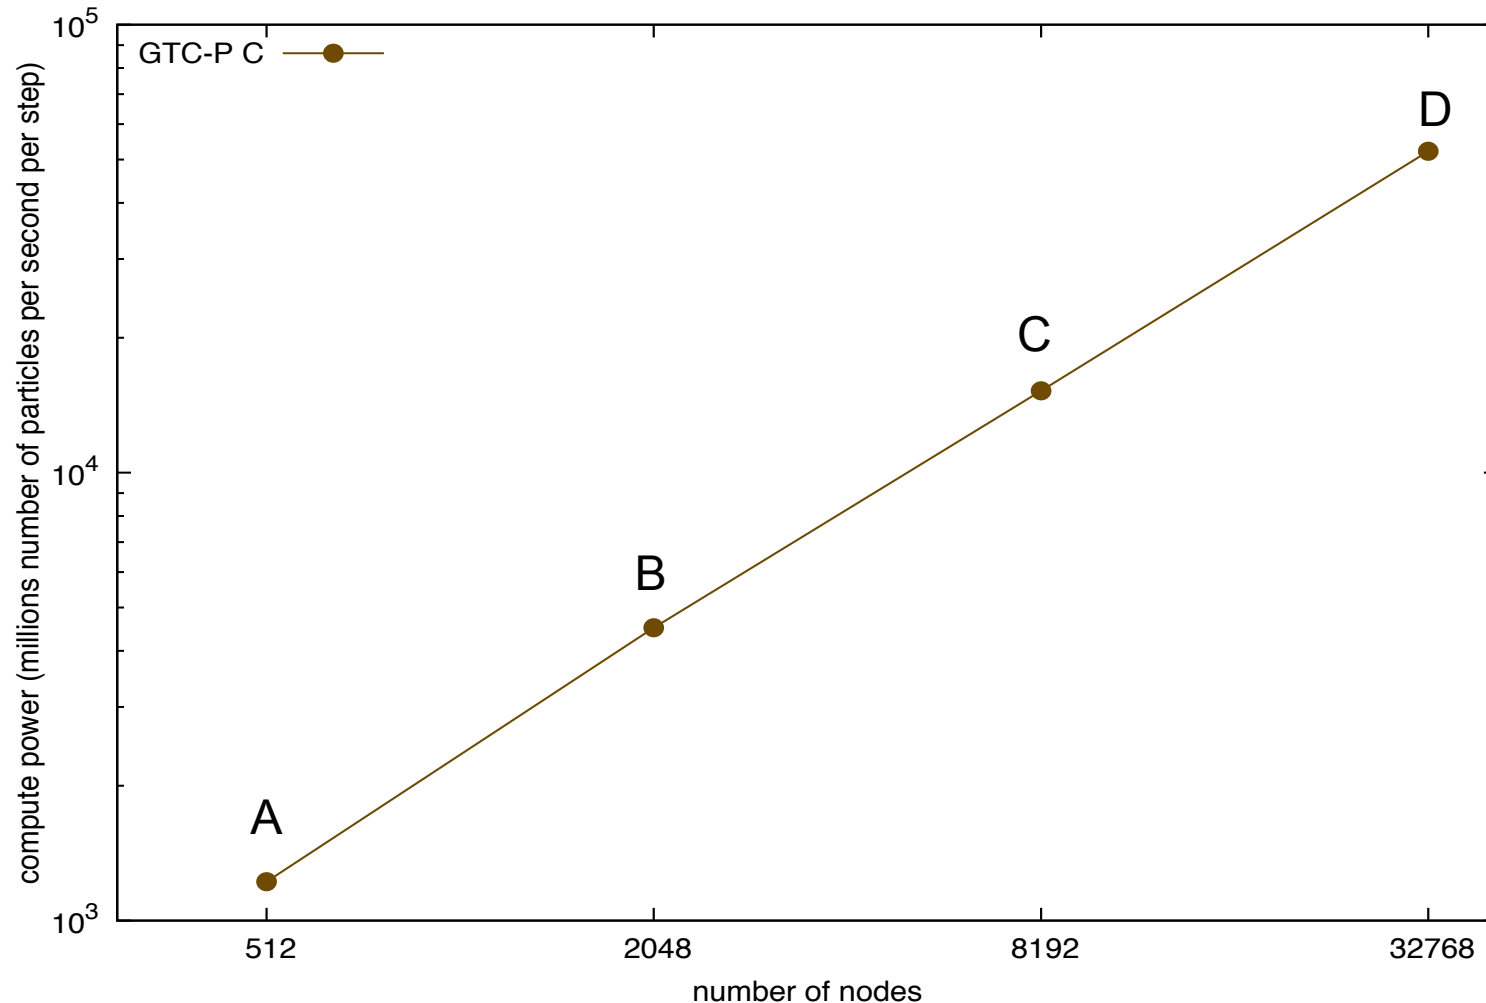
Weak Scaling Study of GTCP-C on Mira



- GTCP-C Titan points beyond 8192 nodes (dashed line) are extrapolated
- 32768 nodes represents 2/3 of Mira BG-Q system

K-Computer Performance: Weak Scaling Results

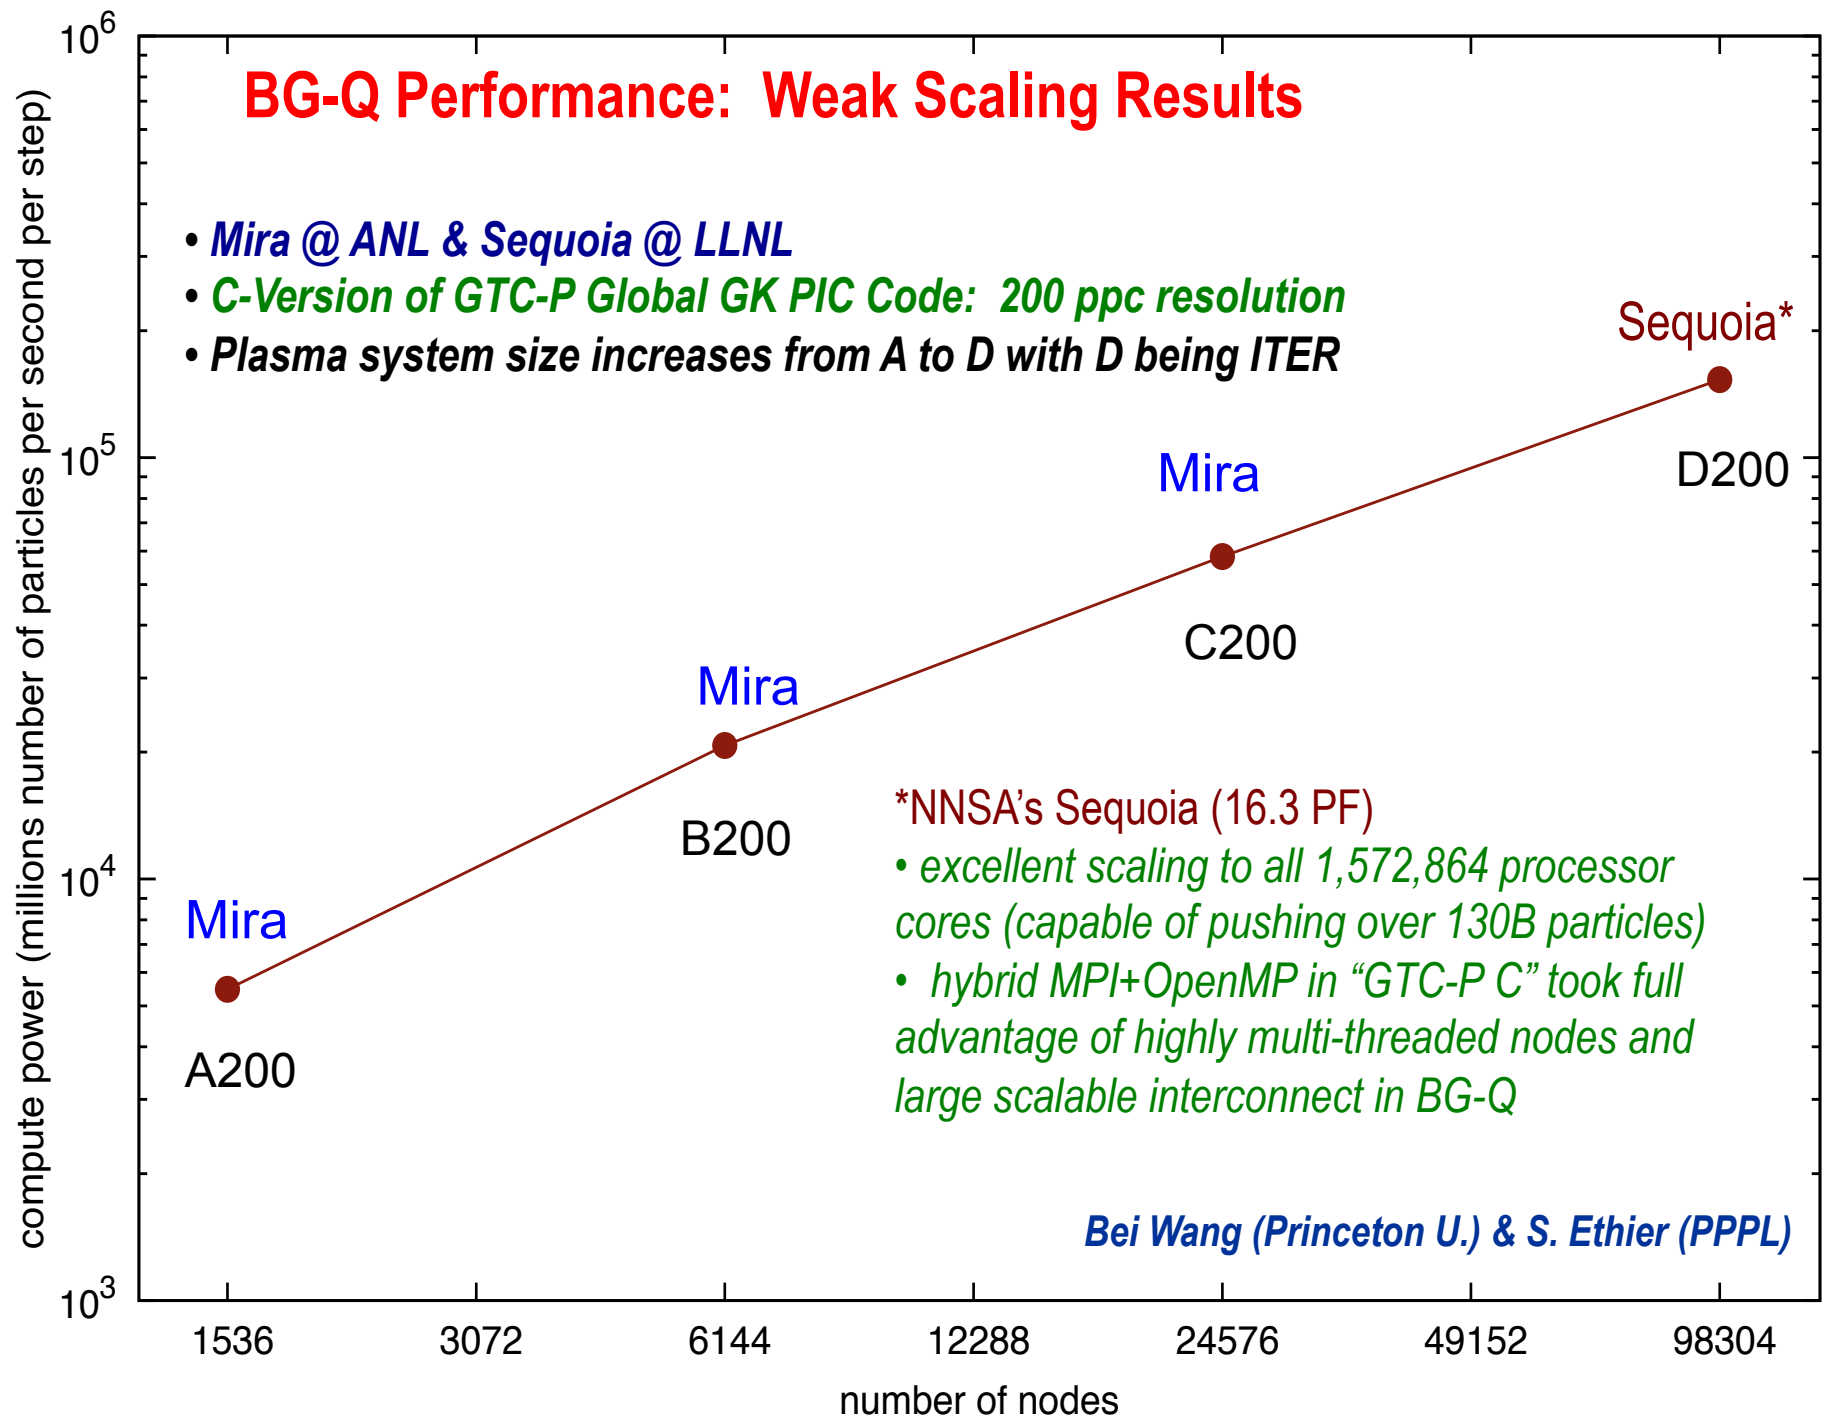
- *Fujitsu-K Computer @ RIKEN AICS, Kobe, Japan*
- *C-Version of GTC-P Global GK PIC Code: 200 ppc resolution*
- *Plasma system size increases from A to D with D being ITER*



Takenori Shimosaka (RIKEN) & Bei Wang (Princeton U.)

BG-Q Performance: Weak Scaling Results

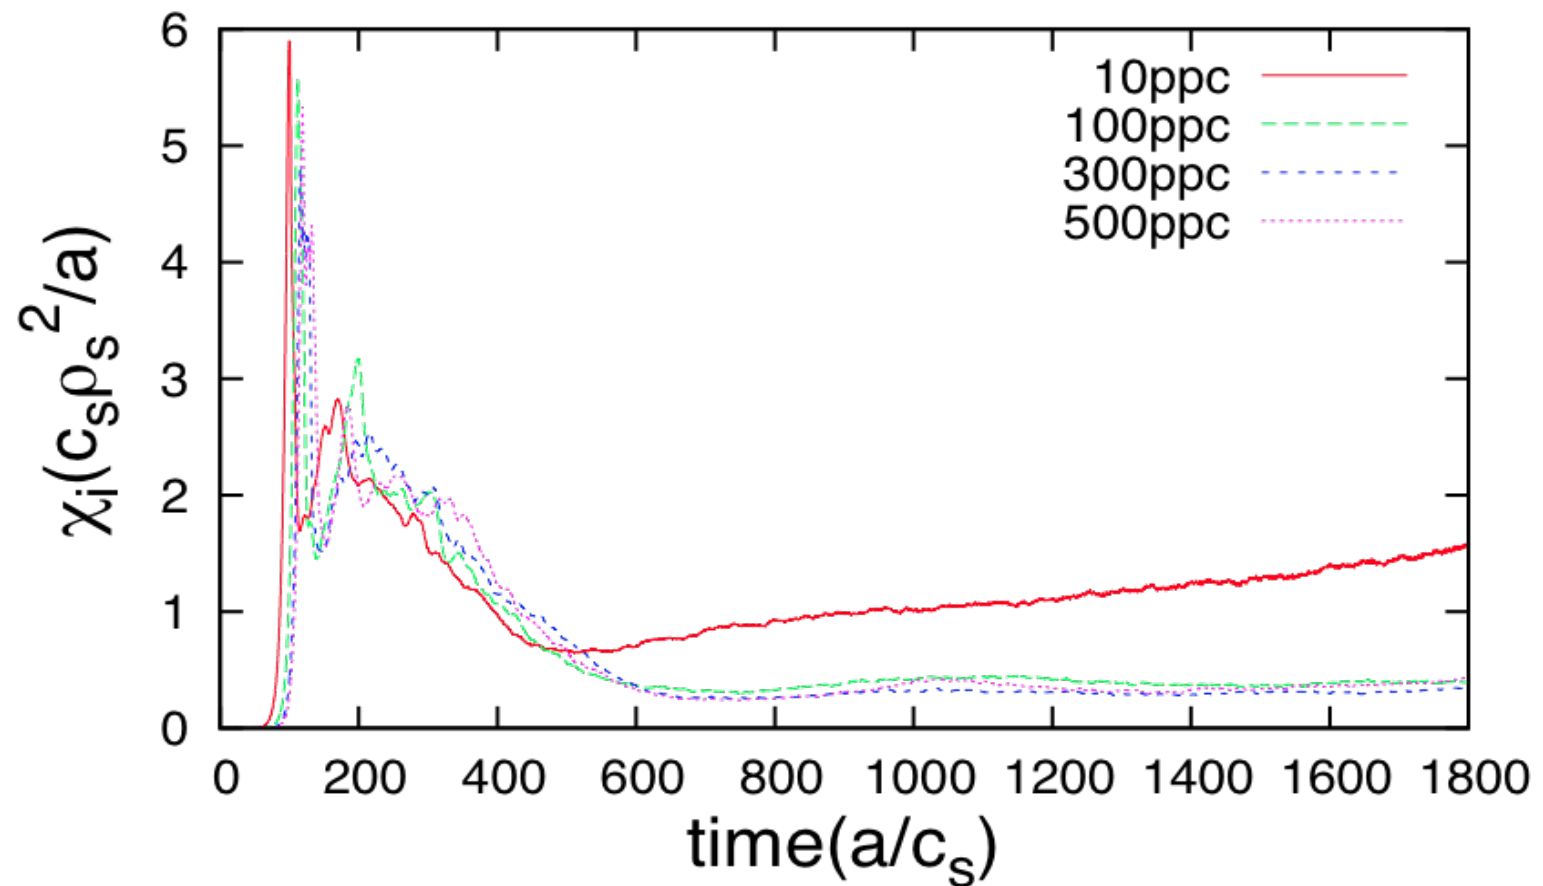
- *Mira @ ANL & Sequoia @ LLNL*
- *C-Version of GTC-P Global GK PIC Code: 200 ppc resolution*
- *Plasma system size increases from A to D with D being ITER*



Particle Resolution (ppc) Convergence Study

GTC-P C Code for ITER (D-size) Case on BG-Q

Time History of Thermal Diffusivity from ITG Instability



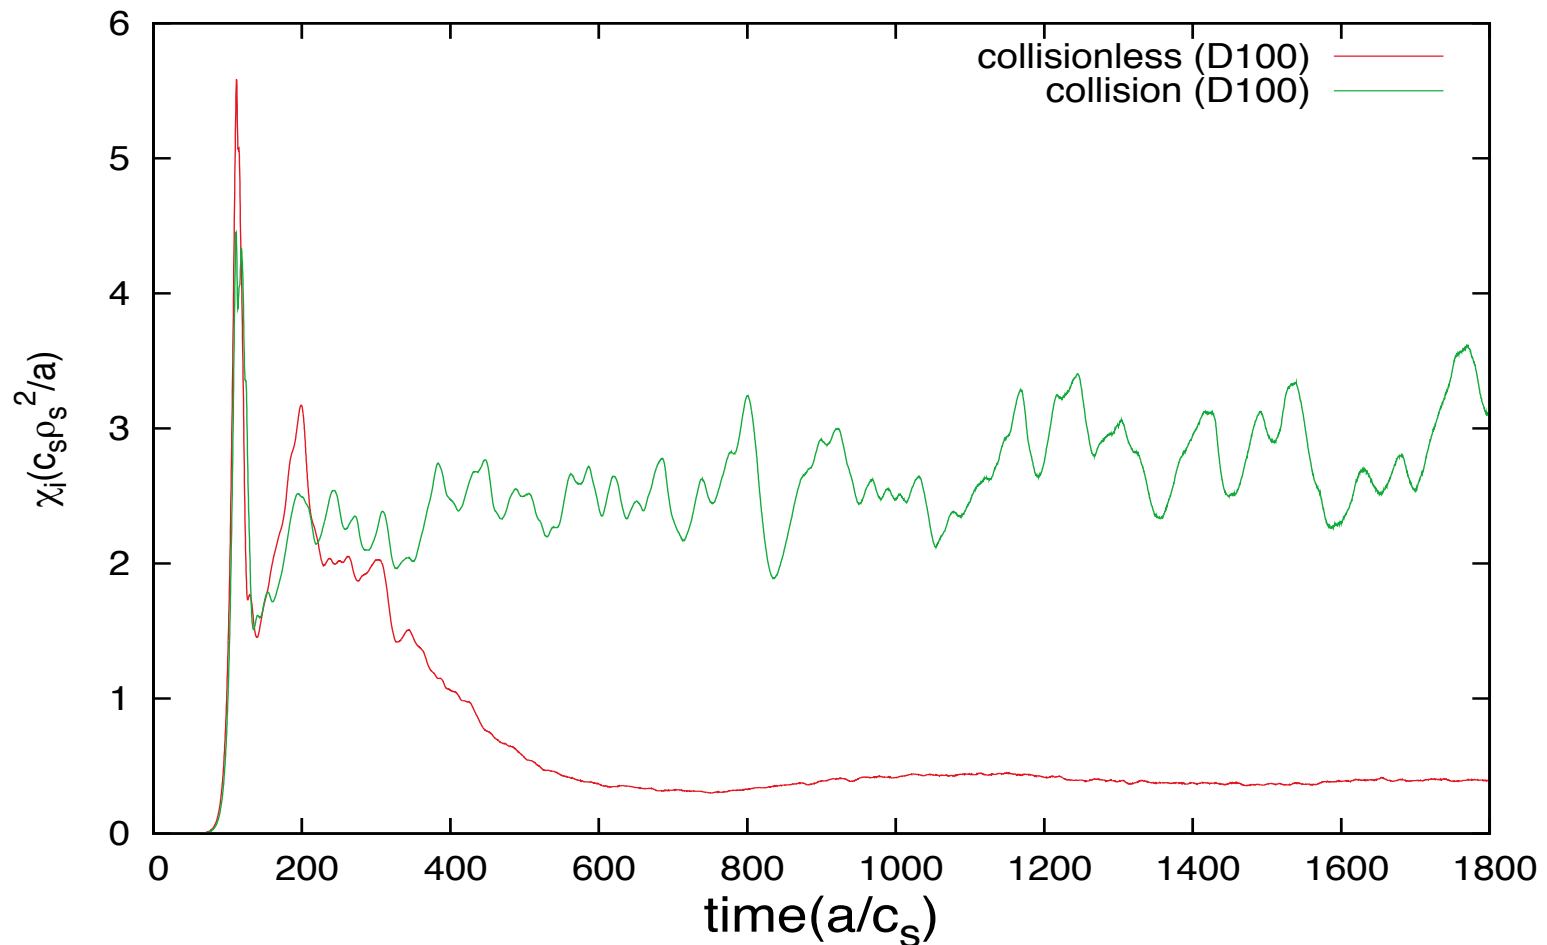
Influence of Collisions on Thermal Diffusivity: Fokker-Planck Model*

GTC-P C Code for ITER (D-size) Case on BG-Q with 100 ppc

*Ref. S. Ethier, W. M. Tang, R. Walkup, and L. Oliner, IBM Journal of R & D, 52 (1-2) 105-115 (2005)

Time History of ITG-driven Thermal Diffusivity (with & without collisions)

➔ *higher heat flux in NL saturated state (as expected) in presence of collisional dissipation*



Summary: Programming Model Challenges in Moving toward Extreme Scales

- **Locality:** Need to improve *data locality* (e.g., by sorting particles according to their positions on grid)
 - due to physical limitations, *moving data between, and even within, modern microchips is more time-consuming than performing computations!*
 - scientific codes often use data structures that are easy to implement quickly but limit flexibility and scalability in the long run
- **Latency:** Further exploration of *highly multi-threaded algorithms* to address memory latency motivated, e.g., by positive results from present studies
- **Flops vs. Memory:** Need to utilize Flops (cheap) to better utilize Memory (limited & expensive to access)
- **Advanced Architectures:** Need more *“demo-apps”* that deploy innovative algorithms within modern science codes on *low memory per core architectures* – (e.g, BG/Q, Fujitsu-K, Titan, Tianhe-1A,)
 - multi-threading within nodes, *maximizes locality while minimizing communications*
 - large future simulations (PIC → very high-resolution (ppc) production runs for long-duration in large-plasma-size scaling studies)

Encouraging performance achieved with “GTC-P C” code on BG/Q (Mira & Sequoia), on Fujitsu-K Computer (Japan), and also CPU part of Titan (OLCF)